(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 2 135 394 B1

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention of the grant of the patent:
15.09.2010 Bulletin 2010/37

(21) Application number: **08718613.6**

(22) Date of filing: **05.03.2008**

(51) Int Cl.:
H04L 12/46 *(2006.01)* ***H04L 12/56*** *(2006.01)*

(86) International application number:
PCT/GB2008/000752

(87) International publication number:
WO 2008/110751 (18.09.2008 Gazette 2008/38)

(54) **MPLS transport network scheme**

MPLS-Transportnetzschema

Schéma de réseau de transport MPLS

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MT NL NO PL PT RO SE SI SK TR

(30) Priority: **14.03.2007 EP 07251059**

(43) Date of publication of application:
23.12.2009 Bulletin 2009/52

(73) Proprietor: **BRITISH TELECOMMUNICATIONS public limited company**
London
EC1A 7AJ (GB)

(72) Inventors:
- **MCGUIRE, Alan**
  Suffolk IP11 7RH (GB)
- **HARRSON, Neil**
  Devon, TG9 6AN (GB)

(74) Representative: **Lofting, Coreena Fiona Anne**
BT Group Legal
Intellectual Property Department
PP C5A, BT Centre
81 Newgate Street
London
EC1A 7AJ (GB)

(56) References cited:
EP-A- 1 134 938

- **BRYANT EDITOR CISCO SYSTEMS S: "Application of PWE3 to MPLS Transport Networks" IETF STANDARD-WORKING-DRAFT, INTERNET ENGINEERING TASK FORCE, IETF, CH, 13 October 2006 (2006-10-13), XP015047504 ISSN: 0000-0004 cited in the application**

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 2 135 394 B1

Description

[0001] The present invention relates to a method of partitioning a multi-protocol label switched (MPLS) header stack into different domains and in particular but not exclusively to a method of sending traffic conforming to one instance of a MPLS transport network protocol over another instance of a MPLS transport network protocol (known in the art as MPLS over MPLS).

[0002] MPLS was originally conceived as a method to improve the packet forwarding performance of Internet Protocol (IP) routers. However, it has subsequently been extended to carry other layer network technologies (like Asynchronous Transfer Mode (ATM), Frame Relay (FR), Plesiochronous Digital Hierarchy (PDH), etc) by the use of Pseudowire (PW) encapsulation techniques. The architecture of the MPLS standard is set out in RFC 3031 entitled "Multiprotocol Label Switching Architecture" submitted by E. Rosen et al in January 2001 to the Internet Engineering Task Force (IETF). Electronic copies of this document are available for download from the URL: www.ietf.org/rfc/rfc3031.txt.

[0003] As is well known in the art, MPLS essentially enables faster routing decisions by preconfiguring "tags" which determine a path between one router and the next. The "tags" are essentially labels carried in short packet header fields which are extracted by switching/forwarding network nodes (known as label switched routers (LSRs)). LSRs are preconfigured to associate certain labels with particular outgoing port(s) and hence traffic containing that label can be routed without a more detailed inspection of the packet header. This avoids the need for hop-by-hop routing decisions to be made on the IP layer network address, instead traffic is sent along a path predetermined by a particular set of labels.

[0004] MPLS label stacking is already known in the art as a means of implementing MPLS tunnelling. To implement MPLS tunnelling, an outer transport label is used to establish a bulk transport Label Switched Path (LSP) (which functions as a tunnel), often between the provider edge devices of a provider's network, and within each bulk LSP inner transport labels are used to identify each traffic flow. Each packet can carry many labelled headers organized as a last-in-first-out stack. In normal forwarding across a MPLS network a LSR processes only the top (i.e. outermost) label. At any LSR, a labelled header can be added to the stack (by the LSR performing a "push" operation) or removed from the stack (by the LSR performing a "pop" operation). The label stacking allows the aggregation of LSPs into a single LSP for a portion of the route, which creates a "tunnel".

[0005] Figure 1A shows a MPLS header. The MPLS header is 32 bits in length and comprises several functional fields. A MPLS stack comprises a plurality of MPLS headers and the position of the MPLS label stack in a frame (or equivalently in a packet) is shown in Figure 1B.

[0006] As shown in Figure 1A, a conventional MPLS 32 bit header comprises a 20 bit label field, a 3 bit EXP (experimental) field, a 1 bit S (stack) field and an 8 bit TTL field. Here the TTL field indicates the time to live of the MPLS packet. The S field is the bottom of stack field which indicates where the MPLS header is in a MPLS stack, S=1 indicates the MPLS header is at the bottom (i.e. innermost position which is closest to the payload) position in a MPLS stack, S=0 if not. Figure 1B shows the how the MPLS label stack entries are positioned above the layer 2 (for example, data link) layer headers but below the network or layer 3 (for example, IP) headers. Other headers may be added, for example, there could also be a PW Control Word header here if the payload data is not IP.

[0007] Figure 1C shows more clearly for a MPLS stack how the top of the label stack is received first by an LSR after the layer 2 (for example, a data link layer) header and the bottom of the label stack appears last and closest to the layer 3 (for example, a network layer) header. This bottom header has the S bit set to 1. All other headers will have the S bit set to 0.

[0008] Figure 2 shows an expanded view of a MPLS Stack comprising three MPLS headers. The top section of Figure 2 shows how the three tiers of MPLS labelled headers are stacked within a packet header. The bottom section of Figure 2 shows how the layer 2 header is followed by two MPLS headers with S=0 and a third MPLS header which appears at the bottom of the MPLS stack has S=1.

[0009] The theoretical number of MPLS label values which could be provided by the 20 bit label value field of the MPLS header is 2^{20} however, in practice, some label values are reserved and have special significance. For further information, see the Request for Comment RFC 3032 entitled "MPLS Label Stack Encoding" submitted to the internet Engineering Task Force standards body (the "IETF") by E. Rosen et al and dated January 2001 which indicates that values 0 to 15 are reserved.

[0010] When a labelled packet is received by an LSR, the label value at the top of the label stack, i.e., closest to the layer 2 header, is first processed to look up forwarding information. This allows the receiving LSR to learn the port to which the packet can be forwarded and/or any operation(s) to be performed on the MPLS stack before forwarding. Examples of operations a LSR performs on a MPLS stack include replacing the top label stack entry with another value and/or popping an entry off the label stack and/or replacing the top label stack entry and/or pushing one or more additional entries onto the label stack. Other operations a LSR can perform include learning the outgoing data link encapsulation and any other information needed to properly forward the packet.

[0011] Depending on the value assigned to the MPLS label field and other fields in the MPLS header stack, the label field can be interpreted in several ways by a receiving LSR. The label field can function, for example, as a forwarding label, as a source label, or as a functional label to indicate an operation that the LSR must perform.

[0012] When used as a forwarding label the label value field of the MPLS header functions as a proxy identifier for the "address" of the LSP destination end-point. In the case where this MPLS header is at the top of a stack of LSP MPLS headers, S=0, if it is however, the sole MPLS header in the stack, then S=1. It is used to forward the MPLS traffic unit in the data-plane towards the destination on a hop-by-hop basis (except when PHP (penultimate-hop-popping) is used, in which case the forwarding header is removed completely on the last hop).

[0013] When used as a source label, the 20 bit label value field functions as a proxy identifier for the "address" of the LSP source end-point. This functional use of a labelled header means it will invariably be the last entry in a stack of MPLS headers and so the S bit will be set to 1.

[0014] When used as a functional label, the label value field identifies a specific action that needs to be taken when received at an LSR. As is known in the art, the reserved values for the label field of the MPLS header specify specific functions which an LSR is to perform on a received packet. For example, label 14 indicates a router OAM alert to pass the contents of the traffic unit to the control/management plane for further action. In such a labelled header the S bit will usually be set to 1 (but this may not always be the case with other functional uses).

[0015] A network architecture has been proposed for transporting a client, MPLS network over a server MPLS network by S. Bryant and is described in the Internet Draft entitled "Application of PWE3 to MPLS Transport Networks" submitted to the Internet Engineering Task Force on 13th October 2006. A copy of this paper can be obtained from https://tools.ietf.org/id/draft-bryant-pwe3-mpls-transport-00.txt. Figure 3A shows how S. Bryant et al have proposed a prior art solution to the problem of how to transport traffic from one operator's MPLS network (MPLS network "A") over another operator's MPLS network (MPLS network "B"). This prior art solution comprises a form of MPLS tunnelling. In Figure 3A, a MPLS Label Switched Path (LSP) is established between Customer Edge nodes CE1 and CE2 which are within the domain of MPLS network "A". This LSP is carried over Ethernet. Ethernet is provisioned between CE1 and PE1, and an Ethernet pseudowire (PW) is provisioned between the Provider Edge equipment PE1 and PE2 to carry the Ethernet traffic from PE1 to PE2, Ethernet is also provisioned between PE2 and CE2. The Ethernet PW is carried over an IP or MPLS PSN. If MPLS is used the tunnel LSP is not configured with PHP.

[0016] The Ethernet transport PW shown in Figure 3A can support multiple service LSP instances, for example a service LSP and a trunk LSP (which can carry aggregations of the service instances). The identification of each service instance within its aggregate is based on labels and service instances can be aggregated onto trunks by pushing labels and de-aggregated by popping labels (here the terms "pushing" and "popping" are used in the sense known in the art).

[0017] The prior art scheme proposed by S. Bryant separates the two MPLS header stacks by inserting an Ethernet header between the two stacks, as is shown schematically in Figure 3B. Thus it is known in the prior art to use an inserted Ethernet layer network header to functionally separate two MPLS layer networks, i.e., to provide a network stack in which MPLS is layered over Ethernet which is in turn layered over MPLS. This prior art solution separates the two MPLS stacks and allows full functional decoupling of the two MPLS networks, e.g. the data/control/management planes of operator A are isolated from the data/control/management planes of operator B. However, it has certain limitations such as, for example, the fact that both networks must support the same buffer technology (i.e., both the server and client MPLS networks must support the same inserted Ethernet technology) and this must go end-end between the CE devices. This is because the PE devices are expecting an Ethernet presentation, which will become, a client PW carried over the middle server PSN network. Another limitation is that inserting a layer 2 technology such as Ethernet into the MPLS stack to enforce a partition results In additional operational overhead, e.g. the management of a degenerate (i.e. 1-hop) Ethernet layer network.

[0018] Accordingly, in the prior art, it is known for the Ethernet layer to extend end-end along the LSP between the customer edge (CE1 and CE2) devices.

[0019] It is also known in the art to support a label switched path across a non-MPLS network. In European Patent Application EP 1 134 938 in the name of Nortel Networks Corporation, Barry Hass describes a system, device and method for supporting a label switched path across a non-MPLS network. This system requires the label stack information of a packet sent over the LSP to be preserved by establishing an IP tunnel across the non-MPLS domain of the LSP which connects a first MPLS domain and a second MPLS domain of the LSP. The packet and label stack information are encapsulated and sent over the IP tunnel and a MPLS identifier is included in the header of the encapsulated packet so that the packet and label stack information may be identified at the second MPLS domain.

[0020] The MPLS transport network scheme of the invention seeks to obviate the limitations of known methods for using MPLS as a transport network, which require a consistent technology to be implemented along a LSP between two client MPLS devices over the transport provider's MPLS network. The invention seeks to provide a method of using MPLS as a transport network comprising interconnecting two Label Switched Routers (LSRs) located in a first (client) MPLS network using the infrastructure of a second (server) MPLS network, i.e., to implement a MPLS-over-MPLS client/server relationship which does not need to insert a different and fixed-type of technology in the MPLS header stack to create a functional separation of the two MPLS networks.

SUMMARY STATEMENTS OF INVENTION

[0021] A first aspect of the invention comprises a method defined by the independent claim 1.

[0022] Another aspect of the invention comprises apparatus defined by independent claim 11.

[0023] Thus a label stack can be delimited by marking an operational boundary so that only a portion of the stack up to a boundary marker (or within a pair of boundary markers) can be operated by one or more nodes at the edge or within an operational domain of the communications system. Examples of operational domains include routing domains.

[0024] In one embodiment, the method enables a multi-protocol label switching (MPLS) transport network scheme to be implemented according to the invention.

[0025] In one embodiment, said step of receiving said packet and processing the header to add a stack entry having a delimiting value is iterated more than once to partition the header stack into additional stack domains.

[0026] In one embodiment, each additional stack domain is partitioned by a stack entry having a different delimiting value.

[0027] In one embodiment, each additional stack domain is partitioned by a stack entry having the same delimiting value.

[0028] In one embodiment, the value assigned to the label field of a header which delimits the header stack entries comprises a reserved value according to the label switching protocol.

[0029] In one embodiment, said stack entry having a delimiting value is indicative that the header stack is to be partitioned at that point such that other nodes in the same routing domain as said edge node are only capable of operating on one more additional headers added to the header stack that were applied by the same routing domain that added the stack delimiting value.

[0030] In one embodiment, said label switching protocol comprises a multi-protocol label switching (MPLS) protocol.

[0031] In this way, one MPLS network is transparently transported over another MPLS network by inserting a special reserved label into the header stack. The hierarchy of MPLS networks comprises a client/server network hierarchy in one embodiment of the invention.

[0032] Another aspect of the invention seeks to provide a method of removing boundary markings in a header stack of a traffic unit said header stack conforming to a label switching protocol and comprising a plurality of headers, each of which comprises a label field to which a label valve is assignable to enable a label switching node to operate on said traffic unit in a network operational domain of a label switching protocol in a communications system comprising a plurality of networks, the method comprising the steps of:

receiving said traffic unit at an edge node of a routing domain in said communications system;

processing the label header stack at the edge node to remove a stack entry having a label field comprising a value which delimits the header stack entries between a stack domain configurable by said node belonging to a first routing domain of said label switching protocol and at least one other stack domain configurable by a node belonging to another routing domain of said communications system.

The header stack is de-partitioned by removing the partition in the header stack. Each part of the label stack reverts back to its disjoint form. In this way, after de-partitioning (i.e., after removing the boundary markers in the stack), the two label stacks are divided either physically or logically. For example, the two label stacks are physically/spatially disjoint when in their parent MPLS networks and/or are logically disjoint when one stack is carried over the other stack.

[0033] In any given data packet, a number of headers are appended according to one or more protocols (for example see Figure 2), for example, there may be layer 2 and/or layer 3 headers present. An example of a header stack conforming to a label switching protocol is a multi-protocol label switching header stack.

[0034] The label field to which a label value is assignable enables a label switching node to perform a particular operation on the packet in the network routing domain, for example, routing, forwarding, etc. Not all labels will have a routing (for example a destination address (DA) proxy forwarding) semantic in some embodiments of the invention. Some labels have a source semantic, some have a local action semantic, etc. Accordingly, the label field enables the label switching node to take the appropriate action on the assumed label semantic. Another aspect of the invention seeks to provide a communications signal comprising a plurality of traffic units, each traffic unit comprising a header and a payload, the header including a header stack comprising a plurality of header fields, all conforming to the same label switching protocol, wherein at least one header field is assigned a value denoting that after that header the stack entries for said header conforming to said label switching protocol are operable on by nodes in a different operational domain of said-label switching protocol to the operational domain of nodes conforming to said label switching protocol which previously operated on the stack entries of the header field.

[0035] The operational domain may comprise a routing domain, or forwarding domain or any other suitable domain of action that a node can perform.

[0036] In some embodiments of the invention, the label field values in the header stack are added by different types of label distribution protocols.

[0037] The MPLS header fields may be contributed by networks associated with different domains of control. In some embodiments of the invention, the stack of MPLS headers all belong to the same network party.

[0038] Another aspect of the invention seeks to provide

a communications system comprising a plurality of multi-protocol label switching communication networks, each said communication network comprising a plurality of network nodes arranged to implement a multi-protocol label switching operational domain, the system comprising: a boundary node between two of said communication networks comprising: a receiver arranged to receive a multi-protocol label switching traffic unit; and one or more processors arranged to process a multi-protocol label switching header stack of said received traffic unit to include a delimiting label field value which indicates an operational boundary of the multi-protocol label switching operational domains of said two communications networks.

[0039] In one embodiment, said node is an ingress node to one of said two communication networks, and wherein said node further comprises one or more processors arranged to: process the header stack after said operational boundary label field value has been assigned to include one or more other multi-protocol label switching header field values enabling one or more operations to be performed on said traffic unit by other nodes within the operational domain of said communication network of said ingress node.

[0040] In one embodiment, said delimiting label field value is processed by another node in said communication network, any label field value assigned to said header stack by a node in the operational domain of the other one of said two communication networks is not operated on.

[0041] In one embodiment, a node receiving said traffic unit in said communication network processes said header stack only up to said delimiting label field value and ignores any label field value assigned to said header stack by a node in the operational domain of the other one of said two communication networks.

[0042] In one embodiment, the type of multi-protocol label switching protocol differs between at least two of said plurality of communications networks.

[0043] The partitioning scheme according to the invention treats the IP addresses in the client MPLS layer network as irrelevant and the client IP addresses do not need to be known in the server layer MPLS network. Instead, the invention carries the client layer MPLS network transparently over the server layer MPLS network.

[0044] The aspects of the invention are as set out above and by the accompanying independent claims and the preferred embodiment are as set out above and by the dependent claims. The aspects and preferred embodiments may be combined in any appropriate manner apparent to those of ordinary skill in the art.

[0045] The preferred embodiments of the invention will now be described with reference to the accompanying drawings which are by way of example only and in which:

Figures 1A to 1C show various views of a conventional MPLS header;

Figure 2 shows expanded views of how MPLS fields are stacked in a MPLS header;

Figure 3A shows a prior art MPLS transport network scheme;

Figure 3B shows how an intermediate technology (Ethernet) separates two separate MPLS transport stacks in the prior art scheme shown in Figure 3A;

Figures 4A shows a communications system within which a MPLS transport scheme can be implemented according to one embodiment of the invention;

Figure 4B shows more detail of a MPLS transport network scheme according to one embodiment of the invention;

Figure 4C shows a partitioned MPLS stack according to one embodiment of the invention; Figures 5A and 5B show partitioned MPLS stacks according to one embodiment of the invention, in which the partitioning MPLS stack entry has S=0;

Figures 6A and 6B show partitioned MPLS stacks according to an alternative embodiment of the invention, in which the partitioning MPLS stack entry has S=1;

Figure 7A shows a communications system for which a recursively stacked MPLS transport scheme according to one embodiment of the invention is provided; and

Figure 7B shows how a MPLS stack is multiply partitioned to implement an embodiment of the invention for the communications system shown in Figure 7A.

[0046] The best mode of the invention currently contemplated by the inventors and the preferred embodiments will now be described. Those of ordinary skill in the art will appreciate that it is obvious that a particular component or step includes certain features whose presence is essential for implementing the invention even if such features have not been explicitly recited (this is for the sake of clarity and brevity in the description and their inclusion is considered implicit unless explicit reference is made to the contrary). All technical terms described herein retain the meaning assigned to them by RFC 3031 and RFC 3032 (as referred to hereinabove) unless a contrary indication is explicitly given.

[0047] The term traffic unit refers to a packet, cell, and frame and /or to any other suitable component of traffic. Each type of traffic unit is used herein as a synecdoche for the other types of traffic units and treated equivalently for the purposes of the invention. For consistency with RFCs 3031 and 3032 which refer to the term MPLS "packet", this term shall be used herein for consistency but the reader who is skilled in the art will understand that refer-

ences to a packet are references to any suitably equivalent type of traffic unit unless the context of the word "packet" clearly indicates otherwise.

[0048] Figure 4A shows a communications system 1 according to the invention which comprises a plurality of networks 2, 3 each comprising a plurality of nodes 4a, 4b, 5a, 5b.

[0049] The nodes 4a, 4b, 5a, 5b of each network are arranged to implement label switching or forwarding on received traffic packets, and in one embodiment comprise label switching routers (LSRs) implemented in accordance with a standard multi-protocol label switched (MPLS) communications protocol.

[0050] According to the invention, within each network 2, 3, at least a plurality of nodes 4a, 4b, 5a, 5b are adapted to interface with at least one other of said plurality of networks. These nodes, which comprises externally facing interfaces with another network, are referred to herein as edge nodes. Depending on the direction of the traffic flow on which the edge node operates the label switching/forwarding operation, the edge node either functions as an ingress node for a traffic flow entering into that particular network or as an egress node for a traffic flow leaving that particular network.

[0051] Each network 2, 3 can have the same or a different type of control and/or management plane technology from the other networks in the communications system. Within each network 2, 3, however, all nodes are within the same operational (for example, routing) domain, i.e. in Figure 4A nodes 4a and 4b belong to one operational (for example, routing) domain, whilst nodes 5a and 5b belong to a different operational (for example, routing) domain. In one embodiment of the invention, two or more (or all) of the plurality of networks forming the communications system can have the same administrative domain. In another embodiment of the invention, two or more (or all) of the plurality of networks forming the communications system have different administrative domains.

[0052] Accordingly, the communications system 1 comprises two or more distinct MPLS networks, each MPLS network comprising at least a different instance of control-plane, and in some embodiments all functional components such as the control-plane, management plane and data plane are independent in each of the MPLS networks. These two or more MPLS networks belong to the same operating party (for example, perhaps to different lines of business for the same operator) in one embodiment of the invention, but in alternative embodiments they belong to different operating parties (for example, different network operators).

[0053] As mentioned hereinabove, in Figure 4A of the accompanying drawings an embodiment of the invention is shown in which the plurality of networks comprises two MPLS networks 2, 3 indicated as MPLS network A and MPLS network B. There are two edge nodes A1 and A2 in MPLS network A and two edge nodes B1 and Bn in MPLS network B, and a label switched path (LSP) 6 denoted by the dashed curved arrow runs from A1 via B1 and Bn to A2. Nodes A1 and A2 are within the operational (for example, routing) domain of network A and nodes B1 and Bn are within the operational (for example, routing) domain of network B.

[0054] The term MPLS node is defined to be consistent with RFC 3031 which refers to a MPLS node as a node capable of running MPLS and which is aware of MPLS control protocols, operates one or more routing protocols, and is capable of forwarding packets based on labels. In Figure 4A, nodes 4a, 4b and 5a, 5b form part of separate MPLS network domains. In accordance with RFC 3031, a MPLS network domain comprises a contiguous set of nodes which perform a MPLS operation, such as routing and forwarding, and which are also in one routing domain.

[0055] For clarity, in each of the two MPLS networks 2, 3, only two edge nodes A1 and A2 (for example, LSRs) are shown residing within MPLS network A and only two edge nodes B1 and Bn (also, for example, LSRs) are shown in Figure 4A. However, those of ordinary skill in the art will appreciate that each MPLS network in any practical deployment is likely to comprise a plurality of nodes, and several may provide core and edge node functionality.

[0056] In one embodiment of the invention, the administrative control of network A is different from that of network B, for example, network A is in the domain of control of a first network operator "A", and the administrative control of network B is in the domain of control of another network operator "B". In alternative embodiments, however, the operator entities "A" and "B" may be the same.

[0057] Consider the scenario where operator A would find it advantageous to send traffic (MPLS packetised data) along an LSP 6 from node A1 via nodes B1 and Bn to another node A2 over the MPLS network of operator B. This would require packets to enter a MPLS network B via the ingress edge node B1 and to leave MPLS network B at egress edge node Bn. As will be apparent to those of ordinary skill in the art, in practice the LSP 6 may transit several nodes in operator B's network between nodes B1 and Bn.

[0058] Conventionally, when MPLS traffic is transmitted between LSRs in a MPLS network, each LSR consults a routing table comprised of MPLS label information which is known as the Label Information Base (LIB). Each LSR performs a look-up operation on the LIB using the ingress interface and ingress label of a received MPLS packet which enables the egress interface and egress MPLS label to be determined, and the LSR then substitutes the outbound label for the incoming label before forwarding the received packet via the egress interface indicated by the LIB. Although each label is only relevant on a particular link between adjacent LSRs, collectively the LSR tables will establish a Label Switched Path over the MPLS network. A MPLS stack comprises a plurality of MPLS labels and only the outermost (last) label is conventionally used for forwarding. The outermost MPLS label is "popped" (i.e., removed) by indicating an appropri-

ate entry in an LSR label table. This enables one or more MPLS labels to be removed at egress LSRs by providing an appropriate indication in the LIB at that LSR. However, for operator A to use operator B's MPLS network as a transport network in a manner which allows the S bit (and indeed all other functions) to be used independently in the two MPLS networks requires some means of preventing the stack entries associated with network A from being operated on by the LSRs associated with network B. This will enable a client/server relationship to be established with a network A MPLS node functioning in a client role and a network B MPLS node functioning in a server role. This allows, as Figure 4A shows, the MPLS node A1 to be peered directly with the MPLS node A2, and for MPLS nodes B1 and Bn to establish a server transport capability for operator A's traffic.

[0059] This requires a partition mechanism to exist in the MPLS stack so that operator B's edge and core MPLS nodes (B1, B2....Bn) have no ability to control the labelling assigned by MPLS nodes (A1, A2) in the client MPLS network of operator A.

[0060] The partition mechanism of the invention allows such a client/server relationship to be established between the two MPLS networks 2, 3 and, referring now briefly to Figure 4C, addresses the problem associated with adding a new MPLS header 8b from network B to a MPLS stack 11 comprising those MPLS headers 8a already applied by network A. The stack partition 9 proposed by the invention acts as a buffer or network delimiter between the two MPLS stack control domains, i.e., it partitions the network stack entries shown in Figure 4C and enables the S bit to be set independently in the MPLS stack entries for each of the two MPLS networks.

A potential benefit of the invention is the provision of a MPLS transport network scheme in which the invention is not prescriptive regarding the type of section layer technology that interconnects the client devices to the server devices. Whilst in one embodiment, the same communications technology supports communications between the client devices and the servers as each end of the LSP, in alternative embodiments, this need not be so.

[0061] In Figure 4B, the layout of the drawing is such that the differences from the prior art scheme shown in Figure 3A are more apparent yet the elements of the communication system are labelled for consistency with those elements shown In Figure 4A. Figure 4B of the accompanying drawings shows how in one embodiment of the invention, the technology supporting communication between the edge nodes A1 and B1 (CE1 and PE1 in Figure 3A) (i.e., the communications protocol) differs from that used between B1 (PE1) and Bn (the provider equipment PE2 in Figure 3A) and this can differ from that used between Bn (PE2) and A2 (the customer equipment CE2 in Figure 3A). Thus in Figure 4B, between A1 and B1 a first communications protocol is used (corresponding to a first type of section layer technology shown as layer 2 #1 in Figure 4B) whereas between Bn and A2 a different section layer communications protocol is used (shown as layer 2 #2 in Figure 4B). In each case, between B1 and Bn no intermediate end-to-end layer network technology exists - the MPLS traffic of network A is not separated by any intermediate technology headers from that of network B. Although referred to as layer 2 in Figures 4A and 4B, the layers 2 #1 and 2 #2 need not always refer to the OSI protocol hierarchy. In some embodiments of the invention, the section layer refers to any appropriate connection-oriented circuit switched (CO-CS), connection-oriented packet switched (CO-PS), or connectionless packet switched (CL-PS) as appropriate depending on the particular embodiment of the invention. For example, the term layer 2 can be used by those of ordinary skill in the art to refer to a CO-PS layer and the term layer 3 can be used to refer to a CL-PS layer.

[0062] In the context of some embodiments of the invention, the term section layer refers to the very bottom of a network stack (i.e., the real or physical layer) where the information (from the higher layer binary signals) need to be encoded (lexicon mapping) into a form suitable for modulation on an electromagnetic wave (which may be guided on a solid physical medium such as copper or optical fibre or comprise a free-space electromagnetic wave propagating as a radio wave for example). As MPLS does not have a section layer specification, it relies on having some other layer network technology (even if a degenerate network such as a point-to-point one hop network) below the MPLS layer. In some embodiments of the invention, these lower layer networks are true networks, in the sense that they create multi-hop switched connections in their own right.

[0063] In Figure 4C a single MPLS stack is partitioned between stack entries whose labels have been assigned values by devices within network A and stack entries whose labels have been assigned values by devices within network B. The header 11 of the MPLS packet shown comprises a layer 2 header component 7, a layer 3 header component 10, and a MPLS header stack 12. In this embodiment, the MPLS header stack 12 comprises two MPLS network headers (8a and 8b) with a network delimiter stack entry 9 providing a functional partition between the sets of stack entries 8a, 8b which belong to the different routing domains of MPLS network A and MPLS network B respectively. The network delimiter stack entry 9 enables the two layer MPLS networks to implement a network hierarchy, for example, by operating in a client/server sense.

[0064] In this way, the invention enables a single MPLS stack to carry MPLS headers conforming to the different values assigned in each network and to keep the MPLS networks functionally independent. Here the terms layer 2 and layer 3 can refer to OSI layers and/or, in some embodiments, layer 2 simply refers to any appropriate connection-oriented packet switched layer technology and layer 3 to any appropriate connection-less packet switched layer technology.

[0065] The invention seeks to define a predetermined value which is assigned to the 20-bit MPLS label field

and which either alone or in conjunction with other MPLS header fields indicates that the MPLS header stack is subdivided into sub-partitions which each contain header values assigned independently by apparatus within a particular MPLS network domain. By ensuring that each MPLS network domain does not attempt to reassign values to the header fields set in another MPLS network domain, the above mentioned limitations in the art as regards MPLS transport networks are mitigated and/or obviated. Advantageously, by partitioning the MPLS header stack into differing network domains it enables recursive use of the MPLS header stack so that multi-depth transport networks can be hierarchically provided such as are shown schematically in Figures 7A and 7B of the accompanying drawings and described later herein below.

[0066] The invention comprises inserting a stack partition value into the label field of a MPLS stack entry which has the functionality of separating the MPLS stack into two or more separate stack control domains, each stack control domain being associated with a separate MPLS network. In this way, the stack partition value functions as a network delimiter. Stack entries whose values have been assigned by the control or management plane of one network operator are separated from those entries assigned values by another network operator by the network delimiter stack entry with the stack partition label value. Thus the network delimiter according to the invention comprises a MPLS stack entry whose label field has been assigned the stack partition value by a server MPLS network domain.

[0067] The network delimiter is inserted by the adaptation function between the client and server MPLS networks, which is implemented by the server layer network. A MPLS packet from the client MPLS network (network A referring back to Figure 4A) is presented to the server MPLS network (network B in Figure 4A) and the ingress node B1 at the edge of the server layer MPLS attaches two new MPLS headers to the MPLS stack. The first MPLS stack entry contains the network delimiter (comprising a MPLS header with the predetermined stack partitioning value) and the second header is a "normal" MPLS header that is used for forwarding across the server layer network. At the egress edge node (Bn as shown in Figure 4A) of the server MPLS network B both of these stack entries are removed and the client MPLS packet is presented to the client MPLS network A. The egress edge node Bn will only be able to "see" the "network delimiter" stack entry after the egress edge node Bn has popped its final forwarding header. The revealed network delimiter label tells the edge node Bn that the next header is a MPLS packet belonging to a different MPLS network.

[0068] As each node along an LSP does not add a new labelled header in normal forwarding, if node B1 applies the two headers 'stack delimiter' + 'normal forwarding', the 'normal forwarding' header gets swapped between B2 and B3, between B3 and B4, etc until the packet is received by node Bn. This means that the stack depth remains at two along the LSP from B1 to Bn, and the 'stack delimiter' header is only observed at B1 (the node adding it) and Bn (the node removing it). In some embodiments, more headers are added to the packet enroute between B1 and Bn; but these create a lower LSP between the nodes adding the header(s), e.g., between some arbitrary intermediate nodes Bk and Bm that lie between B1 and Bn. The intermediate nodes Bk and Bm add/remove these additional header(s) respectively so that they are not present when the packet arrives at Bn for example.

[0069] An example of what is described above is shown in Figures 5A and 5B. In Figures 5A and 5B, examples of two-layer MPLS network header stacks are shown. Figure 5A and 5B show that the presence of the predetermined "partition" value in the label field of a stack entry in conjunction with a S=0 bit indicates that MPLS headers above and below the network delimiter belong to different MPLS networks. Each of these will usually be under different administrative control.

[0070] In Figure 5A, there is only one S=1 stack entry in the MPLS stack, which is consistent with current MPLS standards such as those outlined in RFCs 3031, 3032 etc, and within the partition for network B, the S bit is set to 0 for the label stack entries (#B1 and #B2). Figures 5B, 6A and 6B show alternative embodiments which allow within each stack control domain associated with a particular MPLS network for one S=1 bit to be set.

[0071] In Figure 5B, in each partition of the MPLS stack, the stack entry closest to the layer 3 header end of the MPLS stack is assigned the S=1 bit, and the network delimiter is assigned an intermediate S=0 value. Figures 6A and 6B show an alternative embodiment in which the network delimiter is instead not denoted by the S=0 bit but is associated with a S=1 bit, but different embodiments of the invention may assign different values to the S-bit of the network delimiter.

[0072] Referring again to Figure 4A, a MPLS transport scheme according to one embodiment of the invention will now be described in more detail for a left to right traffic flow (as shown by the arrowed dotted line in the drawing). For a bi-directional service, similar functions will be provided in the reverse direction. To implement the transport scheme according to the invention, operator B sets up (by any mechanism) a standard MPLS LSP, referred to herein as LSP_X, between nodes B1 and Bn. LSP_X would normally have a single MPLS labelled header and the forwarding label would usually (but not necessarily always) be swapped on each hop. Besides containing the forwarding label the other MPLS header fields will be encoded as follows: S=1, EXP=any that is deemed to offer sufficient performance for the purposes of the MPLS client packets from network A that this LSP_X will eventually carry, TTL=any length of time that is sufficient to get from B1 to Bn. In one embodiment, PHP is used on the LSP_X but it is not in the best mode of the invention contemplated by the inventors. This is to ensure that the normal MPLS forwarding header, and specifically the last

hop forwarding label therein, is visible at node Bn. This provides a clear indication of which packets received by node Bn are associated with LSP_X and not any other LSP that terminates at node Bn.

[0073] B1 is configured in any conventional way known to those of ordinary skill in the art to expect to receive MPLS packets from node A1 of operator A on some input port of node B1. The MPLS packets from node A1 can have any allowed values of label, S, EXP, and TTL These MPLS packets are to be carried transparently in a hierarchical (for example, client/server) manner by the MPLS network of operator B using the MPLS transport scheme of this embodiment of the invention. The server section layer that connects node A1 and B1 is any suitable communication technology capable of carrying MPLS packets transparently in a hierarchical (for example, client/server) manner, e.g. Ethernet, known to those of ordinary skill in the art. The 'connection' for that communication technology between A1 and B1 is fully functionally terminated at A1 and B1, i.e. it specifically does not extend across operator B's MPLS network.

[0074] Another server section layer technology 'connection' is also established between the nodes Bn and A2 and satisfies the same requirements as that for the server layer connection between A1 and B1 noted above. Whilst in some embodiments of the invention the server section layer technology between A1 and B1 and that between Bn and A2 is different, in other embodiments the same server section layer technology is provided between A1 and B1 and also between Bn and A2. In general, the manner of establishing server section layer technology connections between A1 and B1 and between Bn and A2 has no impact on the provision of a MPLS transport network scheme according to the invention.

[0075] Referring again to Figure 4A of the drawings, when a MPLS packet from node A1 arrives at node B1, node B1 performs at least the following operations on the MPLS stack of the received packet:

(i) adding a MPLS header comprising the network delimiter to the MPLS packet that came from node A1; and

(ii) adding a further MPLS header after the network delimiter header which is of the type that B1 would normally assign to forward packets from B1 to Bn along LSP_X referred to hereinabove.

[0076] The network delimiter added by B1 contains a label field to which is assigned a Stack Partitioning value. This effectively indicates that that part of the MPLS network stack is terminated and that following the network delimiter are MPLS stack entries belonging to a different MPLS network. Other fields of the network delimiter stack entry are also assigned appropriate values by the ingress node B1 for network B such as the S-bit (e.g. S=0 (recommended), though S=1 could also be used), the EXP field (e.g. EXP=any) and the TTL field (TTL =1 recommended). The field value for TTL (TTL=1 recommended though other TTL values could be used) is assigned to prevent the packet being incorrectly forwarded should it get misdelivered.

[0077] The further MPLS header added by B1 to the MPLS stack is also assigned values in the normal way for MPLS headers in network B, e.g., the S-bit is set to S=0 (recommenced, though in principle S=1 could be used if this is a bottom of stack normal forwarding header in network B between B1 and Bn), the EXP field value is set to any that is deemed to offer sufficient performance for the purposes of the MPLS client packets from network A that this LSP_X will carry, and the TTL field value is set to any value sufficient for the packet to get from B1 to Bn.

[0078] The MPLS packets from node A1 when encapsulated as described hereinabove are then transparently transported across operator B's network (in LSP_X) until they arrive at node Bn. When node Bn removes the normal MPLS forwarding header of LSP_X it will expose the network delimiter stack entry. This indicates that there is a further set of one or more MPLS headers in the MPLS stack after the network delimiter stack entry that belong to a different MPLS network, i.e., in this case the MPLS network of operator A. Node Bn then removes the network delimiter stack entry and, because of prior configuration, knows which port to send this out of node Bn on to reach node A2 in operator A's network.

[0079] One embodiment of the invention uses a label signalling protocol, for example LDP or RSVP-TE, to distribute the labels between nodes (e.g. from node Bn to node B1) and configure the nodes. All labels, including the stack partition label, can be distributed in such a manner within the relevant MPLS network. In other embodiments the label distribution and configuration might be carried out by management and/or OSS techniques. These are conventional techniques already known to those of ordinary skill in the art.

[0080] Figures 7A and 7B of the accompanying drawings show how the invention can be extended to refer to a multi-layer case of MPLS networks comprising a hierarchy of client/server relationships. As shown in Figure 7A, MPLS traffic from the network of operator A uses the MPLS network of operator B as a transport network, who in turn uses the MPLS network of operator C as a transport network, who in turn uses the MPLS network of operator D as a transport network. Figure 7B shows a multi-partitioned MPLS stack in which a plurality of stack partitions #1, #2, #3 are shown. This enables the MPLS stack entries of each network operator to remain concatenated together in the MPLS stack through the data-plane encapsulation technique proposed by the invention for MPLS traffic.

[0081] The MPLS stack partition values are attached as network delimiter stack entries by appropriately configuring the ingress and egress LSRs in each network domain by network management and/or the MPLS signalling using any conventional technique known to those of ordinary skill in the art (such as those described briefly

herein above) with reference to the two layer embodiment of the invention.

[0082] One embodiment of the invention does not force a relationship between the stack partitioning value used in network A and any stack partitioning value used in network B. Nor is any relationship required between any of the stack partitioning values used for network delimiters in embodiments of the invention implemented for a three or more layer MPLS transport network (such as those where the MPLS transport network scheme of the invention is used in a recursive stacked manner as is shown in Figures 7A and 7B (for example, MPLS-over-MPLS-over-MPLS, etc.)).

[0083] In one embodiment of the invention, a globally reserved value (i.e. one of those remaining in the reserved label set 0-15) is assigned to the stack partitioning value which all MPLS networks use. Use of a globally reserved value is not essential however to implement the invention. Any value can be used for the stack partitioning value providing that the MPLS network that assigns that value to an LSP can ensure that it is understood within the context of the LSP that carries it. It may or may not be unique within that MPLS network.

[0084] Many modifications and functional equivalents to the features of the invention described herein will be apparent to those of ordinary skill in the art and unless specifically excluded by the description such modified and functionally equivalent features are intended to be implicitly included in the description, as are any features essential to implement the invention whose omission for the sake of clarity and brevity is apparent to those of ordinary skill in the art.

[0085] The term technology may refer to section layer technology in some embodiments of the invention unless the context of the term indicates otherwise. For example, "communications technology supporting communications between edge nodes" may refer to "section layer" technology in an embodiment of the invention.

[0086] The term "multi-protocol label switched" or "multi-protocol label switching" and the acronym "MPLS" is used herein to refer to standard protocols and implicitly to any variant protocols which have similar header structural constraints, such as, for example, Transport MPLS (T-MPLS, TMPLS is being supported by the International Telecommunications Union (ITU) and more information on TMPLS can be found on the web-site http://www.transport-mpls.com) etc. Accordingly, where appropriate in the description hereinabove, the term MPLS should be construed to include variations and derivative protocols associated with the MPLS protocol such as T-MPLS etc.

[0087] The term traffic unit, packet, cell, and frame are all intended to be synecdoches for each other as they are treated equivalently for the purposes of the invention, although RFCs 3031 and 3032 referred to hereinabove use the term MPLS "packets", which shall be the terms used herein for consistency.

[0088] References to layer 2 and layer 3 may refer herein to OSI layer 2 (data-link) and layer 3 (network) technology and/or to connection-less packet switched and connection-oriented packet-switched (or connection-oriented circuit-switched) communications protocols as is apparent to a person of ordinary skill in the art.

[0089] The text of the abstract is repeated herein-below to form part of the description: a method of partitioning the header stack of a data packet, said header stack conforming to a label switching protocol and comprising a plurality of headers each of which comprises a label field to which a label value is assignable to enable a label switching node to route said packet in a network routing domain of label switching protocol in a communications system comprising a plurality of networks, the method comprising the steps of: receiving said packet at an edge node of a routing domain in said communications system; and processing the label header stack at the edge node to add a new stack entry having a label field comprising a value which delimits the header stack entries between a stack domain configurable by said node belonging to a first routing domain of said label switching protocol and at least one other stack domain configurable by a node belonging to another routing domain of said communications system, whereby a MPLS transport network scheme can be implemented.

Claims

1. A method of partitioning a header stack of a traffic unit, said header stack conforming to a label switching protocol and comprising a plurality of headers each of which comprises a label field to which a label value is assignable to enable a label switching node to operate on said traffic unit in a respective operational domain of said label switching protocol communications network in a communications system comprising a plurality of label switching protocol communications networks, the method comprising of:

   receiving said traffic unit at an edge node of a first operational domain of a said label switching protocol communications network in said communications system;

   processing the label header stack at the edge node to add a new stack entry having a label field comprising a value which marks the boundaries of the header stack entries between a stack domain configurable by said node belonging to said first operational domain of a said label switching protocol communications network and at least one other stack domain configurable by a node belonging to another operational domain of another label switching protocol communications network of said communications system.

2. A method as claimed in claim 1, wherein a said label

switching protocol comprises a multi-protocol label switching protocol.

3. A method as claimed in claim 1, wherein a said label switching protocol comprises a transport multi-protocol label switching protocol.

4. A method as claimed in any one of claims 1 to 3, wherein the traffic unit comprises a ' data packet.

5. A method as claimed in any previous claim, wherein the operational domain comprises a routing domain and wherein said label value is assignable to enable a label switching node to operate on said traffic unit by routing said traffic unit in a respective routing domain of the communications network.

6. A method as claimed in any previous claim, comprising iterating receiving said traffic unit and processing the header to add a stack entry having a delimiting value more than once to mark the boundaries of the header stack entries to partition the header stack into additional stack domains.

7. A method as claimed in claim 6, wherein each additional stack domain is partitioned by a stack entry having a different delimiting value.

8. A method as claimed in claim 6, wherein each additional stack domain is partitioned by a stack entry having the same delimiting value.

9. A method as claimed in any previous claim, wherein the value assigned to the label field of a header which delimits the header stack entries comprises a reserved value according to the label switching protocol.

10. A method as claimed in any previous claim, wherein said stack entry having a delimiting value is indicative that the header stack is to be partitioned at that point such that other nodes in the same routing domain as said edge node are only capable of operating on one more additional headers added to the header stack that were applied by the same routing domain that added the stack delimiting value.

11. Apparatus arranged to partition a header stack of a traffic unit in a communications network, said header stack conforming to a label switching protocol and comprising a plurality of headers each of which comprises a label field to which a label value is assignable to enable a label switching node to operate on said traffic unit in a network operational domain of said label switching protocol in a communications system comprising a plurality of networks, the apparatus comprising:

a receiver arranged to receive said traffic unit at an edge node of an operational domain in said communications system;

one or more processors arranged to process the label header stack at the edge node to add a new stack entry having a label field comprising a value which delimits the header stack entries between a stack domain configurable by said node belonging to said operational domain of said label switching protocol and at least one other stack domain configurable by a node belonging to another operational domain of said communications system.

12. A method of removing boundary markings in a header stack of a traffic unit, said header stack conforming to a label switching protocol and comprising a plurality of headers each of which comprises a label field to which a label value is assignable to enable a label switching node to operate on said traffic unit in a network operational domain of a label switching protocol in a communications system comprising a plurality of networks, the method comprising the steps of:

receiving said traffic unit at an edge node of an operational domain of said communications system;

processing the label header stack at the edge node to remove a stack entry having a label field comprising a value which marks a boundary of the header stack entries between a stack domain configurable by said node belonging to a first operational domain of said label switching protocol and at least one other stack domain configurable by a node belonging to another operational domain of said communications system.

13. A communications signal comprising a plurality of traffic units, each traffic unit comprising a header field and a payload, the header field including a header stack comprising a plurality of header fields conforming to at least one label switching protocol,

wherein

at least one header field is assigned a value denoting that after that header the stack entries for said header conforming to said label switching protocol are operable on by nodes in a different operational domain of said label switching protocol communications system to the operational domain of nodes conforming to said label switching protocol which previously operated on the stack entries of the header field.

14. A communications system comprising a plurality of multi-protocol label switching communication networks, each said communication network comprising a plurality of network nodes arranged to imple-

ment a multi-protocol label switching operational domain, the system comprising:

a boundary node between two of said communication networks comprising:

a receiver arranged to receive a multi-protocol label switching traffic unit; and
one or more processors arranged to process a multi-protocol label switching header stack of said received traffic unit to include a delimiting label field value which indicates an operational boundary of the multi-protocol label switching operational domains of said two communications networks.

15. A system as claimed in claim 14, wherein said node is an ingress node to one of said two communication networks, and wherein said node further comprises one or more processors arranged to:
process the header stack after said operational boundary label field value has been assigned to include one or more other multi-protocol label switching header field values enabling one or more operations to be performed on said traffic unit by other nodes within the operational domain of said communication network of said ingress node.

Patentansprüche

1. Verfahren zum Unterteilen eines Paketkopfstapels einer Verkehrseinheit, wobei der Paketkopfstapel einem Labelvermittlungsprotokoll entspricht und mehrere Paketköpfe umfasst, von denen jeder ein Labelfeld umfasst, dem ein Labelwert zugeteilt werden kann, um einen Labelvermittlungsknoten zu befähigen, an der Verkehrseinheit in einer jeweiligen Betriebsdomäne des Labelvermittlungsprotokoll-Kommunikationsnetzes in einem Kommunikationssystem zu arbeiten, das mehrere Labelvermittlungsprotokoll-Kommunikationsnetze umfasst, mit:

Empfangen der Verkehrseinheit an einem Randknoten einer ersten Betriebsdomäne des Labelvermittlungsprotokoll- Kommunikationsnetzes in dem Kommunikationssystem;
Verarbeiten des Labelpaketkopfstapels an dem Randknoten, um einen neuen Stapeleintrag hinzuzufügen, der ein Labelfeld aufweist, das einen Wert umfasst, der die Grenzen der Paketkopfstapeleinträge zwischen einer Stapeldomäne, die von dem Knoten, der zu der ersten Betriebsdomäne des Labelvermittlungsprotokoll-Kommunikationsnetzes gehört, konfiguriert werden kann, und zumindest einer weiteren Stapeldomäne, die von einem Knoten, der zu einer anderen Betriebsdomäne eines anderen Labelvermittlungsprotokoll-Kommunikationsnetzes des Kommunikationssystems gehört, konfiguriert werden kann, markiert.

2. Verfahren nach Anspruch 1, wobei das Labelvermittlungsprotokoll ein Multiprotokoll-Labelvermittlungsprotokoll umfasst.

3. Verfahren nach Anspruch 1, wobei das Labelvermittlungsprotokoll ein Transport-Multiprotokoll-Labelvermittlungsprotoköll umfasst.

4. Verfahren nach irgendeinem der Ansprüche 1 bis 3, wobei die Verkehrseinheit ein Datenpaket umfasst.

5. Verfahren nach irgendeinem vorhergehenden Anspruch, wobei die Betriebsdomäne eine Routing-Domäne umfasst und wobei der Labelwert zugeteilt werden kann, um einen Labelvermittlungsknoten zu befähigen, an der Verkehrseinheit durch Weiterleiten der Verkehrseinheit in einer jeweiligen Routing-Domäne des Kommunikationsnetzes zu arbeiten.

6. Verfahren nach irgendeinem vorhergehenden Anspruch, mit iterativem Empfangen der Verkehrseinheit und Verarbeiten des Paketkopfs, um einen Stapeleintrag hinzuzufügen, der einen Begrenzungswert mehr als einmal hat, um die Grenzen der Paketkopfstapeleinträge zu markieren, um den Paketkopfstapel in zusätzliche Stapeldomänen zu unterteilen.

7. Verfahren nach Anspruch 6, wobei jede zusätzliche Stapeldomäne durch einen Stapeleintrag mit einem unterschiedlichen Begrenzungswert unterteilt ist.

8. Verfahren nach Anspruch 6, wobei jede zusätzliche Stapeldomäne durch einen Stapeleintrag mit demselben Begrenzungswert unterteilt ist.

9. Verfahren nach irgendeinem vorhergehenden Anspruch, wobei der Wert, der dem Labelfeld eines Paketkopfs zugeteilt ist, der die Paketkopfstapeleinträge begrenzt, einen reservierten Wert nach Maßgabe des Labelvermittlungsprotokolls umfasst.

10. Verfahren nach irgendeinem vorhergehenden Anspruch, wobei der Stapeleintrag, der einen Begrenzungswert hat, angibt, dass der Paketkopfstapel an jenem Punkt so zu unterteilen ist, dass andere Knoten in derselben Routing-Domäne wie der Randknoten nur imstande sind, an einem weiteren zusätzlichen der dem Paketkopfstapel hinzugefügten Paketköpfe zu arbeiten, die von derselben Routing-Domäne, die den Stapelbegrenzungswert hinzufügte, angewendet wurden.

11. Vorrichtung zum Unterteilen eines Paketkopfstapels

einer Verkehrseinheit in einem Kommunikationsnetz, wobei der Paketkopfstapel einem Labelvermittlungsprotokoll entspricht und mehrere Paketköpfe umfasst, von denen jeder ein Labelfeld umfasst, dem ein Labelwert zugeteilt werden kann, um einen Labelvermittlungsknoten zu befähigen, an der Verkehrseinheit in einer Netzbetriebsdomäne des Labelvermittlungsprotokolls in einem Kommunikationssystem zu arbeiten, das mehrere Netze umfasst, mit:

einem Empfänger, der dazu angeordnet ist, die Verkehrseinheit an einem Randknoten einer Betriebsdomäne in dem Kommunikationssystem zu empfangen;
einem oder mehreren Prozessoren, die dazu angeordnet sind, den Labelpaketkopfstapel an dem Randknoten zu verarbeiten, um einen neuen Stapeleintrag hinzuzufügen, der ein Labelfeld aufweist, das einen Wert umfasst, der die Paketkopfstapeleinträge zwischen einer Stapeldomäne, die durch den Knoten, der zu der Betriebsdomäne des Labelvermittlungsprotokolls gehört, konfiguriert werden kann, und zumindest einer anderen Stapeldomäne, die von einem Knoten, der zu einer anderen Betriebsdomäne des Kommunikationssystems gehört, konfiguriert werden kann, begrenzt.

12. Verfahren zum Entfernen von Grenzmarkierungen in einem Paketkopfstapel einer Verkehrseinheit, wobei der Paketkopfstapel einem Labelvermittlungsprotokoll entspricht und mehrere Paketköpfe umfasst, von denen jeder ein Labelfeld umfasst, dem ein Labelwert zugeteilt werden kann, um einen Labelvermittlungsknoten zu befähigen, an der Verkehrseinheit in einer Netzbetriebsdomäne des Labelvermittlungsprotokolls in einem Kommunikationssystem zu arbeiten, das mehrere Netze umfasst, mit den Schritten:

Empfangen des Pakets an einem Randknoten einer Betriebsdomäne des Kommunikationssystems;
Verarbeiten des Labelpaketkopfstapels am Randknoten, um einen Stapeleintrag zu entfernen, der ein Labelfeld aufweist, das einen Wert umfasst, der eine Grenze der Paketkopfstapeleinträge zwischen einer Stapeldomäne, die durch den Knoten, der zu einer ersten Betriebsdomäne des Labelvermittlungsprotokolls gehört, konfiguriert werden kann, und zumindest einer anderen Stapeldomäne, die von einem Knoten, der zu einer anderen Betriebsdomäne des Kommunikationssystems gehört, konfiguriert werden kann, markiert.

13. Kommunikationssignal mit mehreren Verkehrseinheiten, wobei jede Verkehrseinheit ein Paketkopffeld und eine Nutzlast umfasst, wobei das Paketkopffeld einen Paketkopfstapel beinhaltet, der mehrere Paketkopffelder umfasst, die zumindest einem Labelvermittlungsprotokoll entsprechen, wobei
zumindest ein Paketkopffeld einem Wert zugeteilt ist, welcher angibt, dass nach jenem Paketkopf die Stapeleinträge für den Paketkopf, der dem Labelvermittlungsprotokoll entspricht, von Knoten in einer unterschiedlichen Betriebsdomäne des Labelvermittlungsprotokoll-Kommunikationssystems zu der Betriebsdomäne von Knoten, die dem Labelvermittlungsprotokoll entsprechen, welches zuvor an den Stapeleinträgen des Paketkopffelds gearbeitet hat, bearbeitet werden können.

14. Kommunikationssystem mit mehreren Multiprotokoll-Labelvermittlungs-Kommunikationsnetzen, wobei jedes Kommunikationsnetz mehrere Netzknoten umfasst, die dazu angeordnet sind, eine Multiprotokoll-Labelvermittlungs-Betriebsdomäne umzusetzen, mit:

einem Grenzknoten zwischen zweien der Kommunikationsnetze, mit: einem Empfänger, der dazu angeordnet ist, eine Multiprotokoll-Labelvermittlungs-Verkehrseinheit zu empfangen; und
einem oder mehreren Prozessoren, die dazu angeordnet sind,
einen Multiprotokoll-Labelvermittlungs-Paketkopfstapel der empfangenen Verkehrseinheit zu verarbeiten, um einen begrenzenden Labelfeldwert einzuschließen, der eine Betriebsgrenze der Multiprotokoll-Labelvermittlungs-Betriebsdomänen der zwei Kommunikationsnetze angibt.

15. System nach Anspruch 14, wobei der Knoten ein Zugangsknoten zu einem der zwei Kommunikationsnetze ist und wobei der Knoten weiterhin einen oder mehrere Prozessoren umfasst, die dazu angeordnet sind:

den Paketkopfstapel zu verarbeiten, nachdem der Betriebsgrenzen-Labelfeldwert zugeteilt worden ist, um einen oder mehrere weitere Multiprotokoll- Labelvermittlungs- Paketkopffeldwerte einzuschließen, die es ermöglichen, dass ein oder mehrere Vorgänge an der Verkehrseinheit durch andere Knoten innerhalb der Betriebsdomäne des Kommunikationsnetzes des Zugangsknotens durchgeführt werden.

Revendications

1. Procédé de partitionnement de la pile d'en-têtes

d'une unité de trafic, ladite pile d'en-têtes se conformant à un protocole de commutation d'étiquette et comprenant une pluralité d'en-têtes dont chacun comprend un champ d'étiquette auquel une valeur d'étiquette est attribuable pour permettre à un noeud de commutation d'étiquette d'effectuer des opérations sur ladite unité de trafic dans un domaine opérationnel respectif dudit réseau de communication à protocole de commutation d'étiquette dans un système de communication comprenant une pluralité de réseaux à protocole de commutation d'étiquette, le procédé comprenant les étapes consistant à :

recevoir ladite unité de trafic au niveau d'un noeud de bordure d'un premier domaine opérationnel dans ledit système de communication à protocole de commutation d'étiquette ;
traiter la pile d'en-têtes d'étiquette au niveau du noeud de bordure pour ajouter un nouvel élément à la pile ayant un champ d'étiquette comprenant une valeur qui délimite des éléments de pile d'en-têtes entre un domaine de pile configurable par ledit noeud appartenant audit premier domaine opérationnel dudit réseau de communication à protocole de commutation d'étiquette et au moins un autre domaine de pile configurable par un noeud appartenant à un autre domaine opérationnel d'un autre réseau de communication à protocole de commutation d'étiquette dudit système de communication.

2. Procédé selon la revendication 1, dans lequel un dit protocole de commutation d'étiquette comprend un protocole de commutation d'étiquette multi-protocole.

3. Procédé selon la revendication 1, dans lequel un dit protocole de commutation d'étiquette comprend un protocole de commutation d'étiquette multi-protocole de transport.

4. Procédé selon l'une quelconque des revendications 1 à 3, dans lequel l'unité de trafic comprend un paquet de données.

5. Procédé selon l'une quelconque des revendications précédentes, dans lequel le domaine opérationnel comprend un domaine de routage et dans lequel ladite valeur d'étiquette est susceptible d'être affectée pour permettre à un noeud de commutation d'étiquette d'effectuer des opérations sur ladite unité de trafic par le routage de ladite unité de trafic dans un domaine de routage respectif du réseau de communication.

6. Procédé selon l'une quelconque des revendications précédentes, comprenant l'itération de la réception de ladite unité de trafic et le traitement de l'en-tête afin d'ajouter plus d'une fois un élément de la pile présentant une valeur délimitante afin de marquer les frontières des éléments de la pile d'en-têtes afin de partitionner la pile d'entêtes en domaines de pile supplémentaires.

7. Procédé selon la revendication 6, dans lequel chaque domaine de pile supplémentaire est partitionné au moyen d'un élément de la pile présentant une valeur de délimitation différente.

8. Procédé selon la revendication 6, dans lequel chaque domaine de pile supplémentaire est partitionné au moyen d'un élément de pile présentant la même valeur de délimitation.

9. Procédé selon l'une quelconque des revendications précédentes, dans lequel la valeur affectée au champ d'étiquette d'un en-tête délimitant les éléments de la pile d'en-têtes comprend une valeur réservée en correspondance au protocole de commutation d'étiquette.

10. Procédé selon l'une quelconque des revendications précédentes, dans lequel ledit élément de pile présentant une valeur délimitante est indicatif du fait que la pile d'en-têtes doit être partitionnée à ce point de telle sorte que d'autres noeuds dans le même domaine de routage que ledit noeud de bordure sont seulement capables d'effectuer des opérations sur un en-tête additionnel supplémentaire ajouté à la pile d'en-têtes qui a été appliqué par le même domaine de routage qui a ajouté la valeur délimitant la pile.

11. Dispositif agencé pour partitionner une pile d'en-têtes d'une unité de trafic dans un réseau de communication, ladite pile d'en-têtes se conformant à un protocole de commutation d'étiquettes et comprenant une pluralité d'en-têtes dont chacun comprend un champ d'étiquette auquel une valeur d'étiquette est susceptible d'être affectée pour permettre à un noeud de commutation d'étiquette d'effectuer des opérations sur ladite unité de trafic dans un domaine opérationnel de réseau dudit protocole de commutation d'étiquette dans un système de communication comprenant une pluralité de réseaux, le dispositif comprenant :

un récepteur agencé pour recevoir ladite unité de trafic au niveau d'un noeud de bordure d'un domaine opérationnel dans ledit système de communication ;
un ou plusieurs processeurs agencés pour traiter la pile d'en-têtes d'étiquette au niveau du noeud de bordure afin d'ajouter un nouvel élément de pile ayant un champ d'étiquette qui comprend une valeur délimitant les éléments de la pile d'en-têtes entre un domaine de pile sus-

ceptible d'être configuré par ledit noeud appartenant audit domaine opérationnel dudit protocole à commutation d'étiquette, et au moins un autre domaine de pile susceptible d'être configuré par un noeud appartenant à un autre domaine opérationnel dudit système de communication.

12. Procédé pour enlever les marquages de frontière dans la pile d'en-têtes d'une unité de trafic, ladite pile d'en-têtes se conformant à un protocole de commutation d'étiquette et comprenant une pluralité d'en-têtes dont chacun comprend un champ d'étiquette auquel une valeur d'étiquette est attribuable pour permettre à un noeud de commutation d'étiquette d'effectuer des opérations sur ladite unité de trafic dans un domaine opérationnel à protocole de commutation d'étiquette dans un système de communication comprenant une pluralité de réseaux, le procédé comprenant les étapes consistant à :

recevoir ladite unité de trafic au niveau d'un noeud de bordure d'un domaine opérationnel dudit système de communication ;
traiter la pile d'en-têtes d'étiquette au niveau du noeud de bordure pour supprimer un élément de la pile ayant un champ d'étiquette comprenant une valeur qui délimite les éléments de pile d'en-têtes entre un domaine de pile configurable par ledit noeud appartenant à un premier domaine opérationnel dudit protocole de commutation d'étiquette et au moins un autre domaine de pile configurable par un noeud appartenant à un autre domaine opérationnel dudit système de communication.

13. Un signal de communication comprenant une pluralité d'unités de trafic, chaque unité de trafic comprenant un champ d'en-tête ainsi qu'une charge utile, le champ d'en-tête comprenant une pile d'en-têtes qui comprend une pluralité de champs d'en-tête se conformant à au moins un protocole de commutation d'étiquette,
dans lequel
à au moins un champ d'en-tête est affectée une valeur qui signifie qu'après cet en-tête, les éléments de pile pour ladite en-tête se conformant audit protocole de commutation d'étiquette sont susceptibles d'être l'objet d'opérations de la part de noeuds se trouvant dans un domaine opérationnel différents dudit système de communication à protocole de commutation d'étiquette des domaines opérationnels de noeuds qui se conforment audit protocole de commutation d'étiquette qui précédemment effectuaient des opérations sur les éléments de pile du champ d'en-tête.

14. Système de communication comprenant une pluralité de réseaux de communication à commutation d'étiquette multi-protocole, chacun parmi lesdits réseaux de communication comprenant une pluralité de noeuds de réseau agencés pour la mise en oeuvre d'un domaine opérationnel à commutation d'étiquette multi-protocole, le système comprenant :

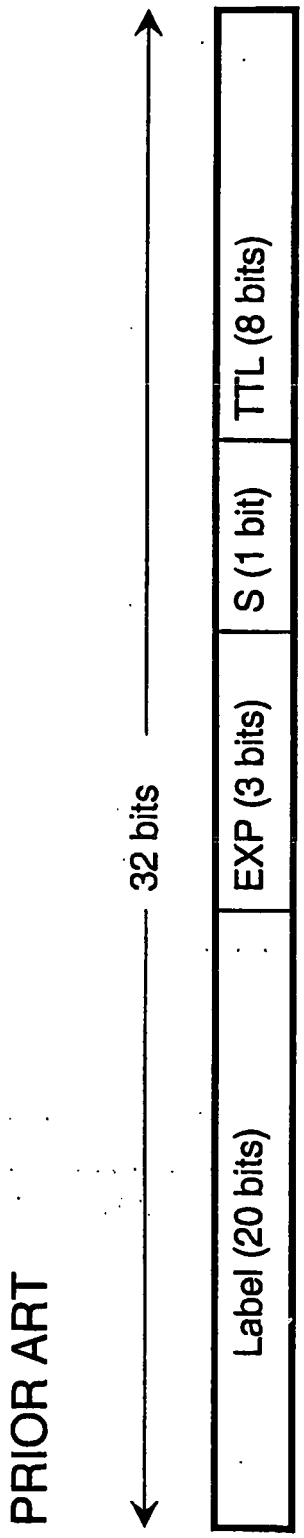
un noeud de frontière entre deux parmi lesdits réseaux de communication comprenant :

un récepteur agencé pour recevoir une unité de trafic à commutation d'étiquette multi-protocole ; et
un ou plusieurs processeurs agencés pour traiter une pile d'en-têtes de commutation d'étiquette multi-protocole de ladite unité de trafic reçue afin d'inclure une valeur de champ d'étiquette délimitante indiquant une frontière opérationnelle des domaines opérationnels à commutation des étiquettes multi-protocole desdits deux réseaux de communication.

15. Système selon la revendication 14, dans lequel ledit noeud est un noeud d'entrée à l'un parmi lesdits deux réseaux de communication, et dans lequel chaque noeud comprend en outre un ou plusieurs processeurs agencé(s) pour :

traiter la pile d'en-têtes après l'affectation de ladite valeur de champ d'étiquette de frontière opérationnelle afin d'inclure une ou plusieurs autres valeurs de champ d'en-tête de commutation d'étiquette multi-protocole permettant la réalisation d'une ou de plusieurs opérations sur ladite unité de trafic par d'autres noeuds à l'intérieur du domaine opérationnel dudit réseau de communication dudit noeud d'entrée.

PRIOR ART



MPLS Header Format

FIGURE 1A

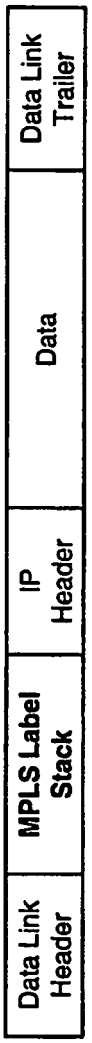


FIGURE 1B

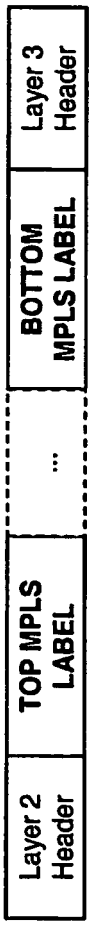


FIGURE 1C

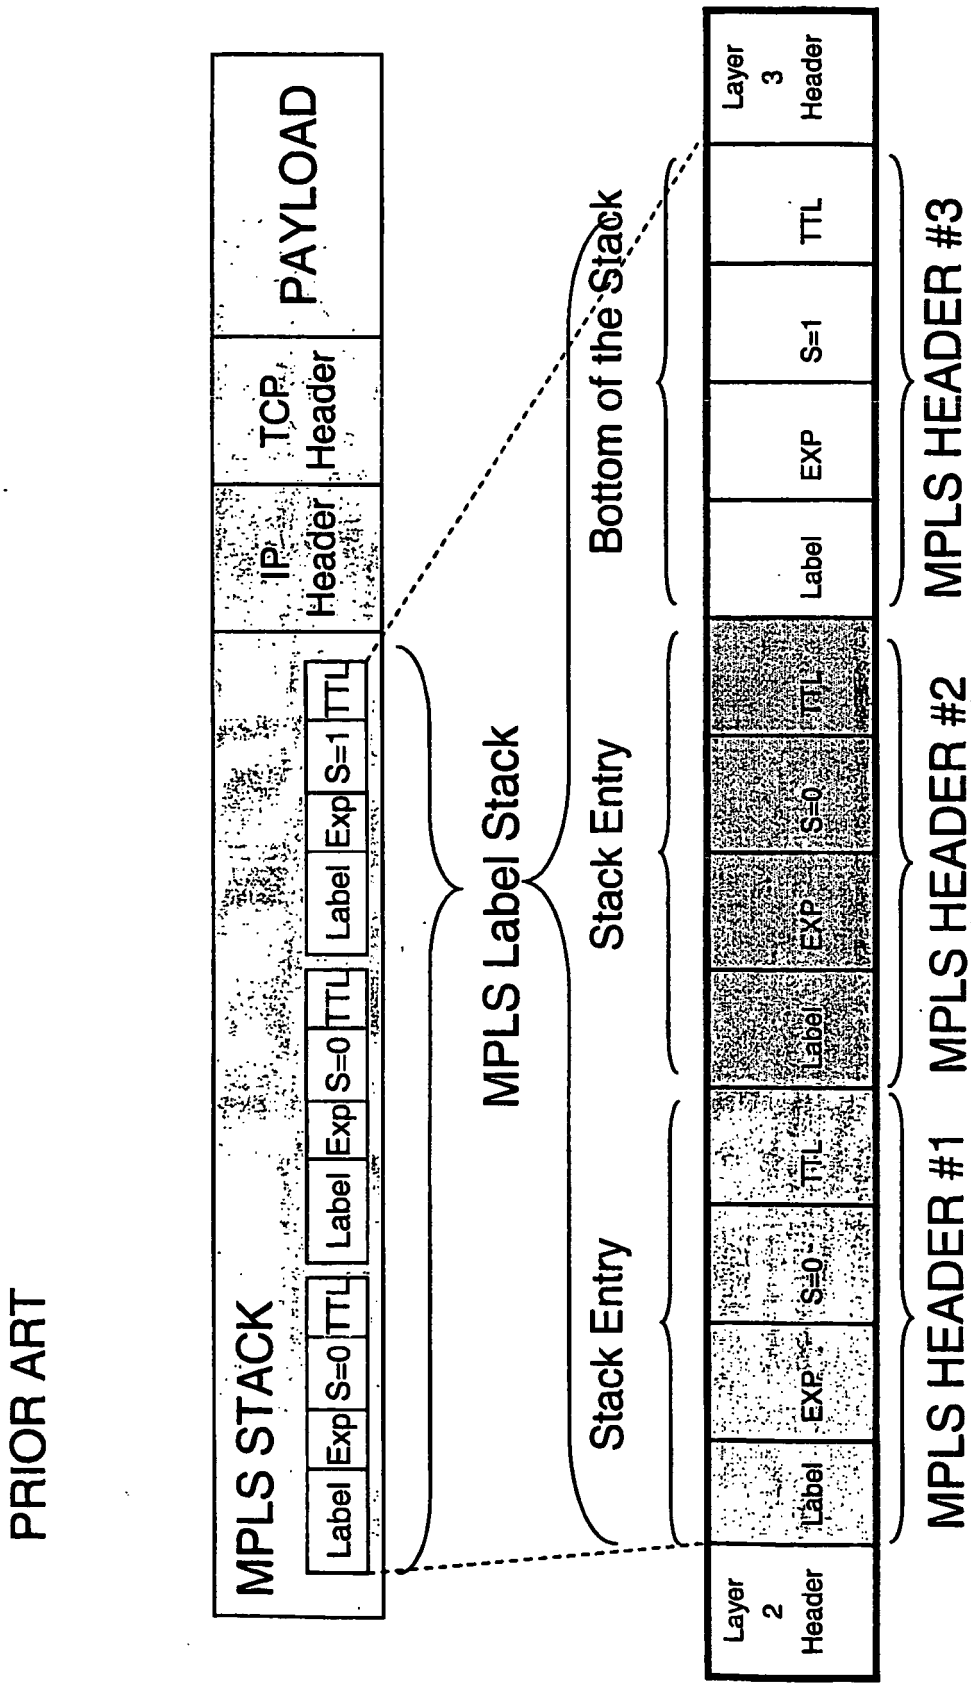

PRIOR ART
MPLS STACK
Label
Exp
S=0
TTL
Label
Exp
S=0
TTL
Label
Exp
S=1
TTL
IP Header
TCP Header
PAYLOAD
MPLS Label Stack
Stack Entry
Stack Entry
Bottom of the Stack
Layer 2 Header
Label
EXP
S=0
TTL
Label
EXP
S=0
TTL
Label
EXP
S=1
TTL
Layer 3 Header
MPLS HEADER #1
MPLS HEADER #2
MPLS HEADER #3

FIGURE 2

PRIOR ART

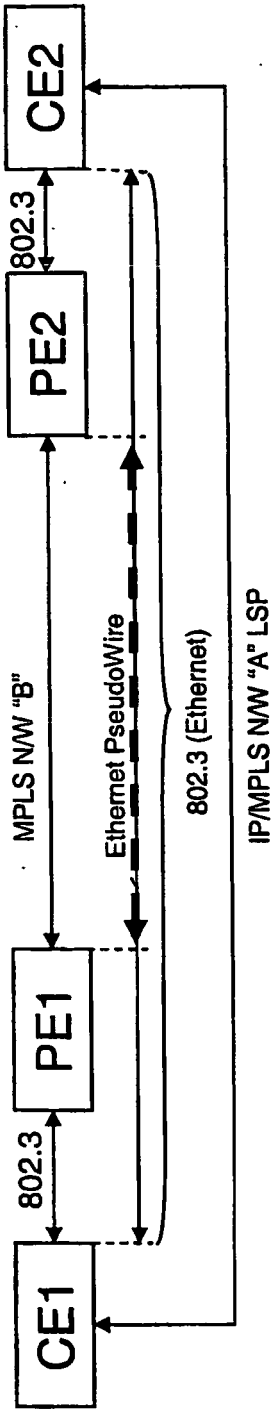


FIGURE 3A



FIGURE 3B

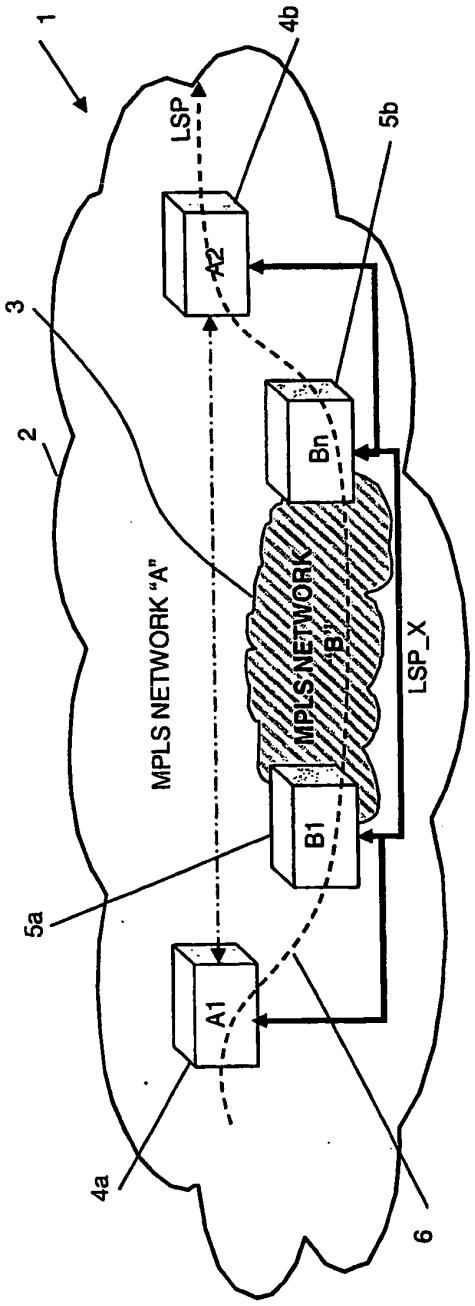


FIGURE 4A

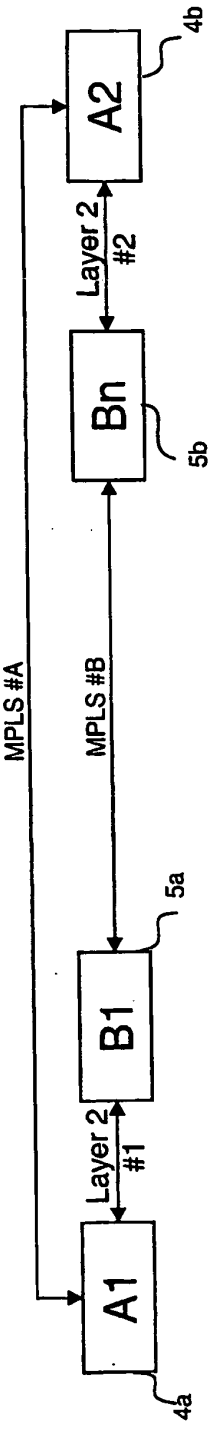


FIGURE 4B

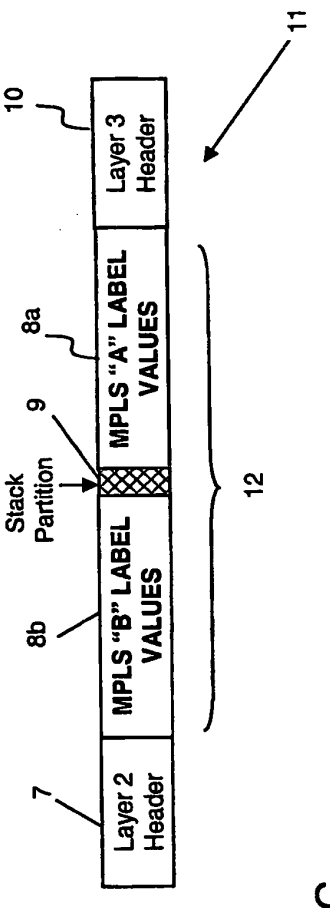


FIGURE 4C

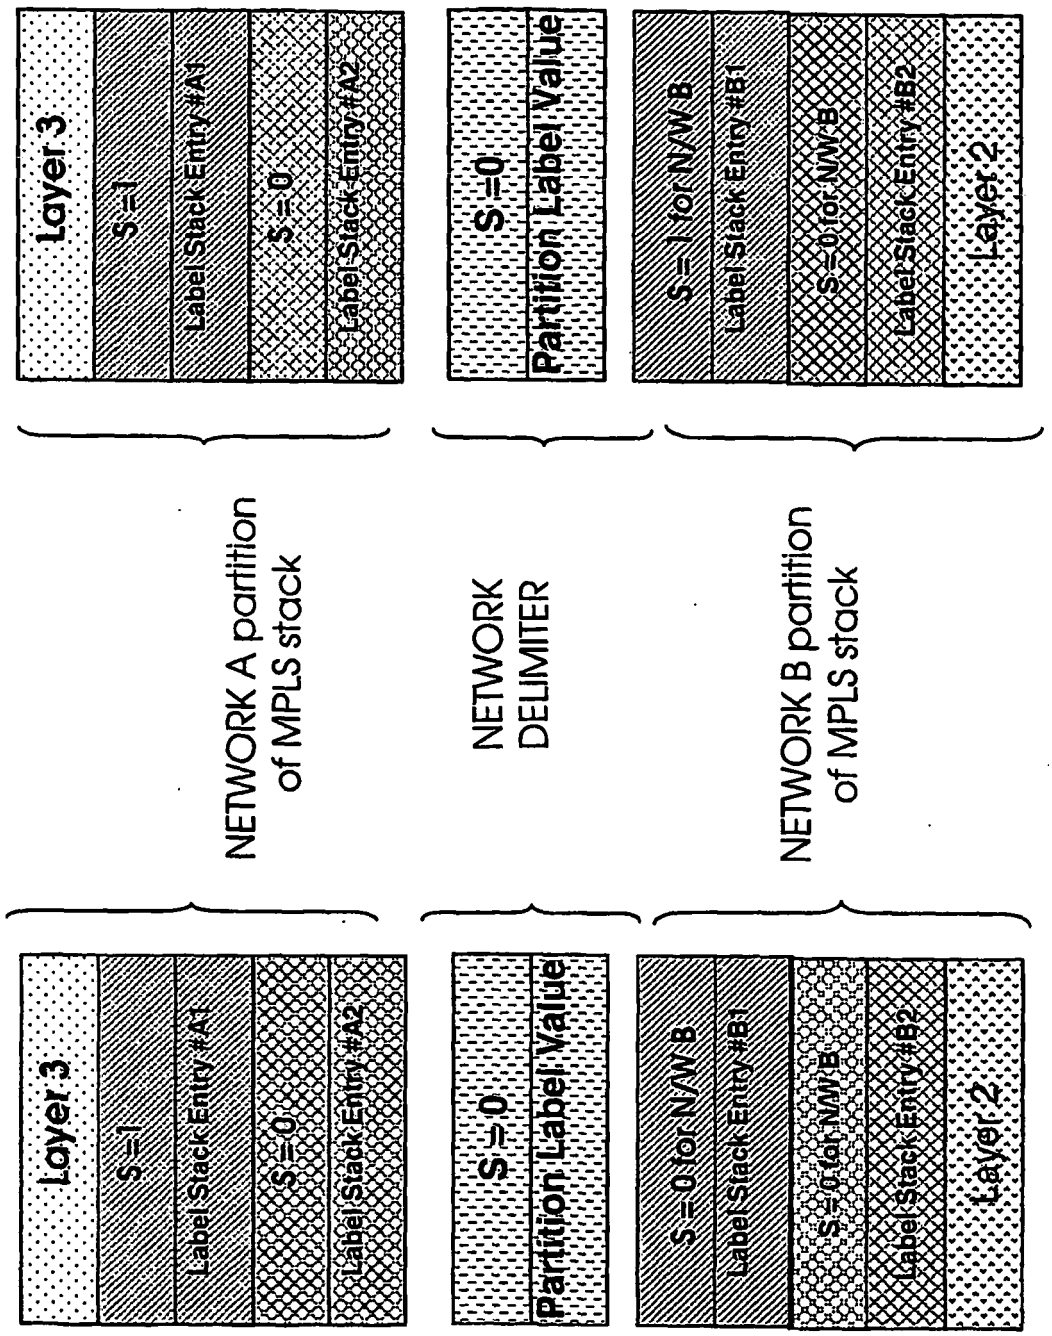


FIGURE 5A

FIGURE 5B

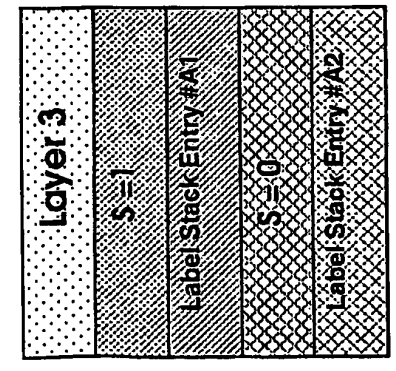


FIGURE 6A

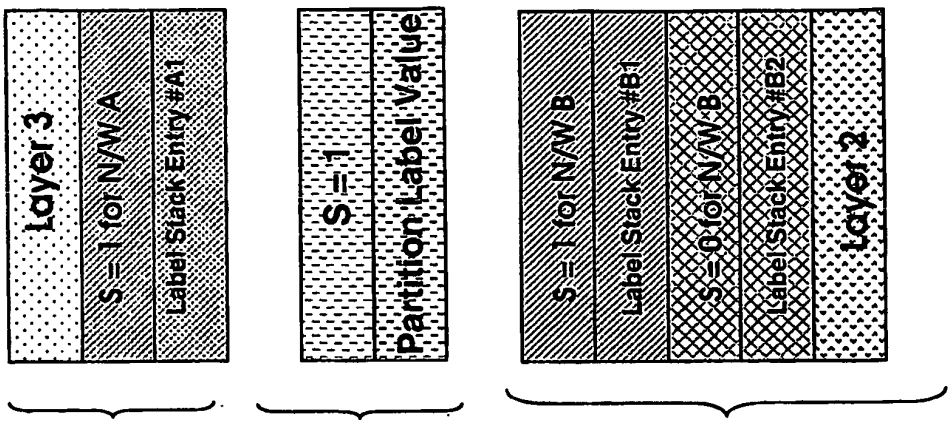


FIGURE 6B

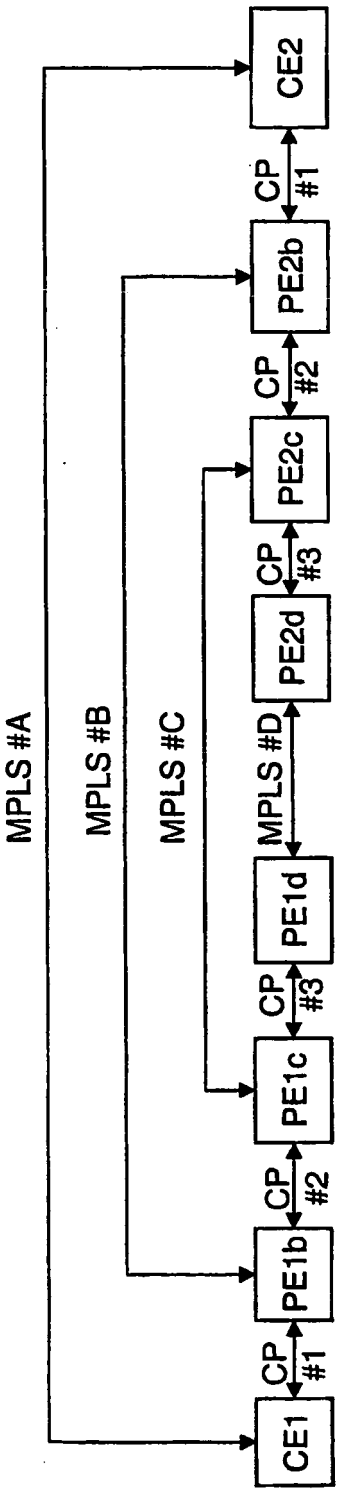


FIGURE 7A

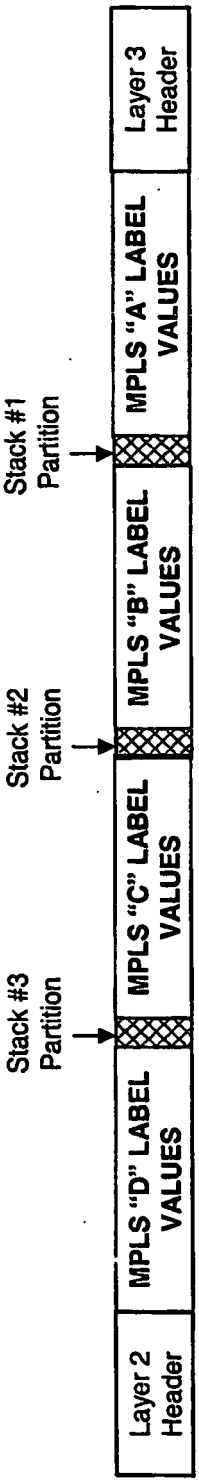


FIGURE 7B

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- EP 1134938 A **[0019]**